Towards CNN Map Compression for camera relocalisation

Luis Contreras[1] and Walterio Mayol-Cuevas[2]

arXiv:1703.00845v1 [cs.CV] 2 Mar 2017

***Abstract*— This paper presents a study on the use of Convolutional Neural Networks for camera relocalisation and its application to map compression. We follow state of the art visual relocalisation results and evaluate response to different data inputs – namely, depth, grayscale, RGB, spatial position and combinations of these. We use a CNN map representation and introduce the notion of CNN map compression by using a smaller CNN architecture. We evaluate our proposal in a series of publicly available datasets. This formulation allows us to improve relocalisation accuracy by increasing the number of training trajectories while maintaining a constant-size CNN.**

I. INTRODUCTION

Following our work on point cloud compression mapping via feature filtering in [1] and [2], we aim to generate compact map representations useful for camera relocalisation via compact CNNs. To ask what is the minimal map representation that enables later use is a meaningful question that underpins many applications for moving agents.

In this work, we specifically explore a neural network architecture tested for the relocalisation task. We study the response of such architecture to different inputs – e.g. color and depth images –, and the relocalisation performance of pre-trained neural networks in different tasks.

Biologically inspired visual models have been proposed for a while [3], [4]. How humans improve learning after multiple training of the same view and how they filter useful information have also been an active field of study. One widely accepted theory of the human visual system suggests that a number of brain layers sequentially interact from the signal stimulus to the abstract concept [5]. Under this paradigm, the first layers – connected directly to the input signal – are a series of specialized filters that extract very specific features, while deeper layers infer more complex information by combining these features.

Finally, overfitting a neural network by excessive training with the same dataset is a well known issue; rather, here we study how the accuracy improves by revisiting the same area several times introducing new views to the dataset.

This paper is organized as follows. In Section II we discuss work related to convolutional neural networks and camera pose. Our CNN for relocalisation is then introduced in Section III, where we describe its architecture. Then, Section IV introduces the notion of CNN map representation and compression. Experimental results are presented in Section V; finally, we outline our discussion and conclusions.

*This work was partially supported by CONACYT and the Secretaria de Educacion Publica, Mexico

Department of Computer Science, University of Bristol, United Kingdom
[1]cslact@bristol.ac.uk
[2]wmayol@cs.bris.ac.uk

II. RELATED WORK

Even though neural networks are not a novel concept, due to the increase in computational power, their popularity has grown in recent years [6] [7].

Related to map compression, dimensionality reduction through neural networks was first discussed in [8]. In [9] an evaluation to up-to-date data encoding algorithms for object recognition was presented, and it was extended in [10] to introduce the use of Convolutional Neural Networks for the same task.

[11] introduced the idea of egomotion in CNN training by concatenating the output of two parallel neural networks with two different views of the same image; at the end, this architecture learns valuable features independent of the point of view.

In [12], Jarrett et al. concluded that sophisticated architectures compensate for lack of training. Garg et al. [13] explore this idea for single view depth estimation where they present a stereopsis based auto-encoder that uses few instances on the KITTI dataset. Then, [14], [15], and [16] continued studying the use of elaborated CNN architectures for depth estimation.

Moving from depth to pose estimation was a logical step. One of the first 6D camera pose regressors was presented in [17] via a general regression NN (GRNN) with synthetic poses. More recently, PoseNet is presented in [18], where they regress the camera pose using a CNN model. This idea is also explored in [19] for image matching via training a CNN for frame interpolation through video sequences.

III. THE RELOCALISATION CNN

We develop a Convolutional Neural Network, or CNN, to address the camera relocalization problem. A CNN can be considered as a filter bank where the filters' weights are such as they minimize the error between an expected output and the system response to a given input. In Figure 1 we show the elements from one layer to the next in a typical CNN architecture – a more detailed CNN implementation can be found in specialized works such as [20] and [21]. For a given input I and a series of k filters f_k, it is generated an output $\hat{I}_k = I * f_k$; the filters f_k can be initialized randomly or pre-trained in a different task.

We highlight the direct relationship between the input channels and the filters' depth, because we will work with different n-dimensional inputs. We evaluate the performance in a number of dataset both RGB and RGBD – more specifically, we tested tensors with the following information per element: gray and RGB values, depth distance, 3D spatial position, RGB+depth, and RGB+spatial position; some typical inputs are shown in Figure 2. The original input is

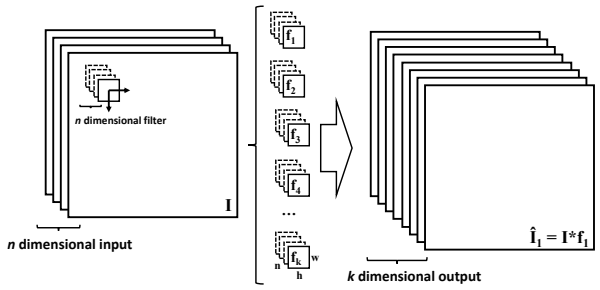

Fig. 1: Convolutional Neural Network (CNN) elements. It consist of an input I, a series filters f_k, and its associated output $\hat{I}_k$. The filter depth depend on the number of input channels.

pre-processed by cropping the central area and resizing it, generating 224x224 arrays.

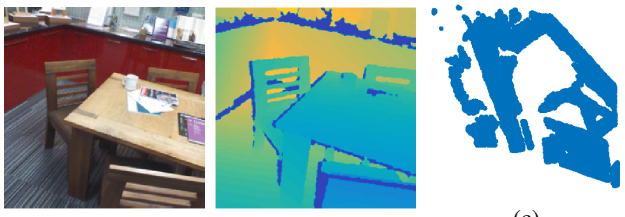

(a) (b) (c) Relocalisation

Fig. 2: Different data input samples: a) color image, b) depth map, and c) 3D point cloud.

Our study is based in the PoseNet model as described in [18]. The expected output is a 6 DOF pose $p = [x, q]$, where x is the spatial position and q is the orientation in quaternion form. We use PoseNet loss function:

$$loss(I) = \|\hat{x} - x\|_2 + \beta \left\| \hat{q} - \frac{q}{\|q\|} \right\|_2$$

For ease, instead of the GoogLeNet arquitecture [22], we use a fast convolutional neural network (CNN-F) implementation as in [10] – an architecture with 8 layers: 5 convolutional and three fully-connected. We introduce a couple of changes, as follows: the dimension in the first layer depends on the input *n*; in addition, the final fully-connected layer size changes to the pose vector length. Table I details our network.

TABLE I: Fast Convolutional Neural Netwirk (CNN-F) architecture, as in [10]. The dimension *n* in *conv1* depends on the input data; the output in *full8* is the pose vector.

conv1	conv2	conv3	conv4
11x11xnx64	5x5x64x256	3x3x256x256	3x3x256x256
conv5	full6	full7	full8
3x3x256x256	4096	4096	7

We tested the CNN-F architecture's performance with the St Marys Church sequence, a large scale outdoor scene [23], as shown in Figure 3. We obtained a relocalization mean error of 8.54 meters, that is in the same magnitude order as that reported in [18] [2.65 m in a modified GoogLeNet architecture with 23 layers].

(a) (b) (c) (d)

Fig. 3: Typical views from the PoseNet dataset [18].

In Figure 4 it can be observed our implementation's output. We use 100 epochs and barely vary the hyperparameters, because at this point the main goal was to design a system that allowed us to perform relocalisation using a CNN. Therefore, we associate the difference in precision compared with that in PoseNet to the mentioned lack of training and the use of a smaller CNN; however, the general performance is maintained.

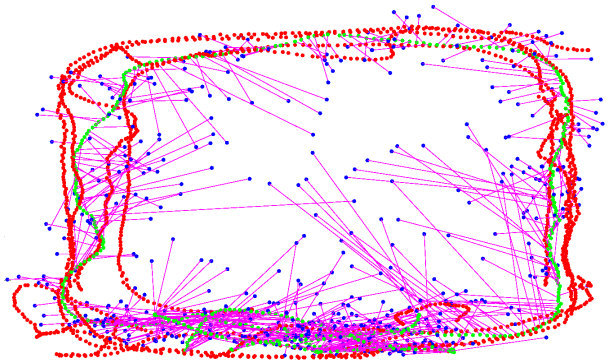

Fig. 4: Relocalisation output applying a CNN-F in the St Marys Church sequence [23]. The red line is the training sequences and the green one the test sequence. Blue points are the output of the system.

IV. CNN Map Representation

From a human observer point of view, it is common to think of metric or topological relationships among elements in space to build maps; for this reason, metric and topological are two well known map representations (such as [24], [25], and [26]). However, other less intuitive map representation have been proposed – e.g. [27] defines a map as a collection of images and uses image batch matching to find the current position in the map.

Overall it can be argued that the map representation need not conform to a *single* representation type, and that the task and other constraints can lead to different manners in which a map can be represented. Ultimately, for a robotic agent, maps are likely built to be explored or more generally re-explored. It is thus that we highlight once more that relocalisation is a good measure of map effectiveness. In this context the actual map representation used is less relevant as long as it improves relocalisation.

In this work, we represent a map as a regression function $\hat{p} = cnn(I)$ where an element of the population is an input I and its associates output p (e.g. a RGB image and the 6DOF camera pose). The parameters in the regressor cnn are optimized from a population sample (training data); the more representative the sample, the more accurate the model [28].

Finally, we introduce the notion of map compression using CNN map representation. We first define a sample as a collection of elements (I, p) that defines a sensor's traveled trajectory. From a series of samples, we define a regressor $\hat{p} = cnn(I)$ that minimizes the error $|p - \hat{p}|$. This regressor, once defined, is of constant size, and should improve its performance while increasing the population sample (i.e. the number of training trajectories). A compact map representation is then posed as the problem of finding an optimal CNN architecture that keeps minimum relocalisation error values.

V. EXPERIMENTS AND RESULTS

A. CNN for camera relocalisation

We first evaluate this implementation with the TUM's long household and office sequence [29] – a texture and structure rich scene, with 21.5m in 87.09s (2585 frames), and a validation data set with 2676 extra frames as can be seen in Figure 5.

(a) (b) (c) (d)

Fig. 5: Some scenes in the TUM [29] long office sequence.

Learning the external parameters in a CNN is time consuming because there is not really an optimal approach to this task; for this reason, the use of generic image representations has been proposed such as [30], where the models trained in one task can be used in another.

We tested the CNN-F with pre-trained wights optimized for object detection in the ImageNet dataset [31], as well as with random weights. The response to a RGB input can be seen in Figure 6 – where the red line indicates the training sequence and the green one the test sequence – with a relocalisation mean error of 0.572 meters using pre-trained weights, and 0.867 meters when random weights were used.

We associate this operation difference to the pre-training process. Figure 7 visualize the filters after 250 epochs; more structured filters – even though they were trained for a different task – showed better relocalisation performance.

Regarding the other sensor inputs, as pre-trained CNN is more common for RGB than depth or spatial data, in this experiment we tested randomly initialized CNNs to make a fairer comparison among them. In Table II we can see the relocalisation mean error for all inputs, where RGB information seems to outperform the other.

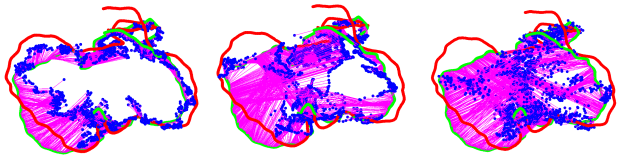

(a) (b) (c)

Fig. 6: Relocalisation performance in the TUM sequence [29]: a) RGB input with a pre-trained CNN, b) RGB input and a randomly initialized CNN, and c) 3D point cloud input. Red line indicates the training trajectory while the green line is the testing one. Blue points are the output of the neural network.

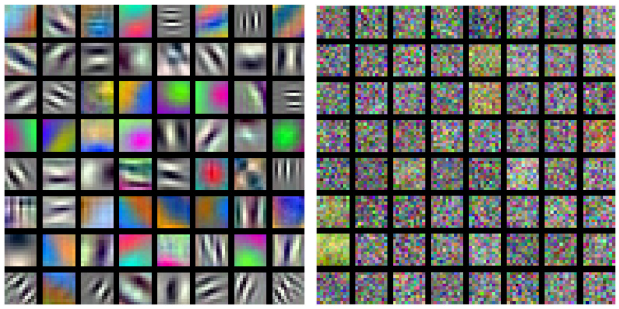

(a) (b)

Fig. 7: First layer filters visualisation for a) pre-trained weights after 250 finetunning training epochs and b) randomly initialized filters after training.

TABLE II: CNN-F relocalisation mean error [in meters] for different inputs after 1000 epochs in the long office sequence from the TUM dataset [29].

	Relocalisation mean error [m]
Depth	1.926 ± 0.595
Gray	0.936 ± 0.814
Point Cloud	1.191 ± 0.960
RGB	0.877 ± 0.820
Pre-trained RGB	**0.572 ± 0.492**
RGB+Depth	1.002 ± 0.817
RGB+Point Cloud	2.542 ± 0.456

We noticed that, when combine information layers are used, the performance decreases. It might be due to the difference in the input nature (color, depth, and spatial position). One way to solve it can be using parallel networks and then average the output, as in [11].

Overfitting – i.e., a parameter saturation in simple model – also affects the performance. In Figure 10 we observe that even though the RGB and point cloud training and validation processes behave similarly, in the test set the differences are notorious. To overcome this issue, we can use a more complex model or pre-trained filters in bigger datasets – the use of 3D primitives to retrieve information useful for image understanding was explored in [33], while [34] learned generative models of 3D structures from volumetric data.

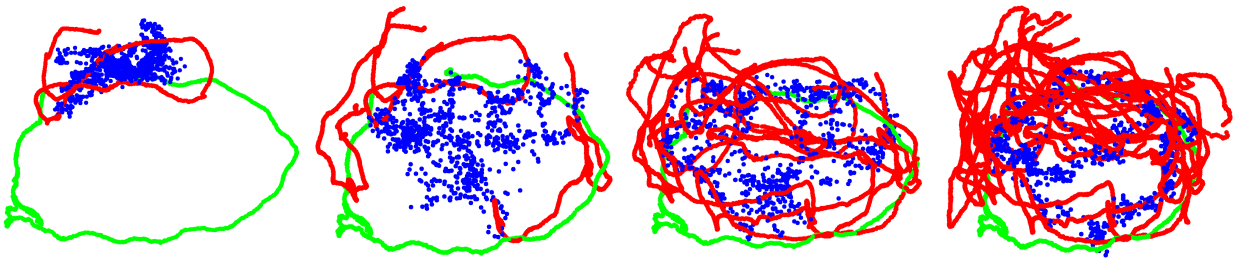

(a) (b) (c) (d)

Fig. 8: Relocalisation performance in from the Red Kitchen sequence in 7-Scenes dataset [32] after training with a) one, b) three, c) seven, and d) eleven sequences. Red lines indicate the training sequences and the green line is the test one; blue points are the output of the system.

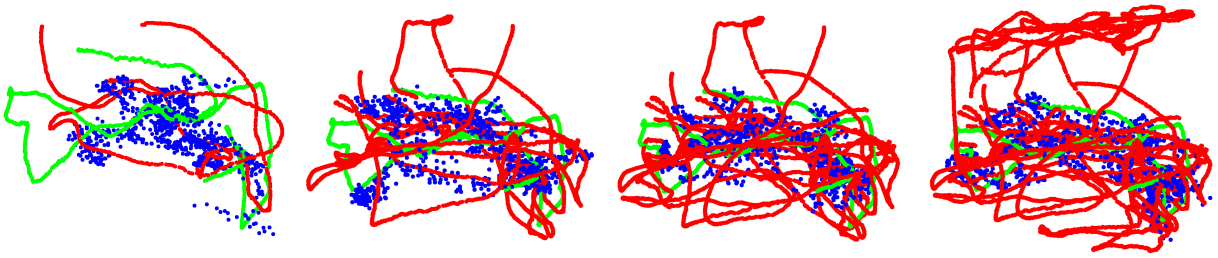

(a) (b) (c) (d)

Fig. 9: Relocalisation performance in the 7-Scenes' Office sequence [32] after training with a) one, b) four, c) six, and d) nine sequences. Colors are as in Figure 8.

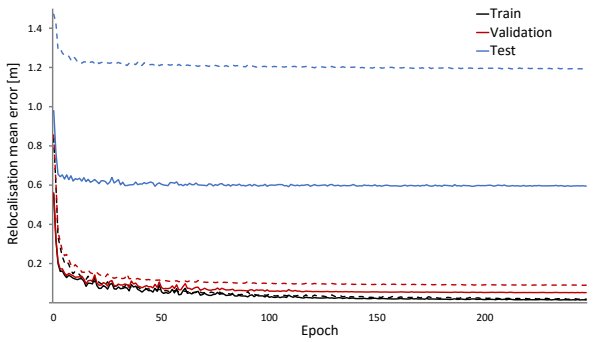

Fig. 10: Relocalisation mean error at different inputs. A continuous line corresponds to an RGB input while a dotted line corresponds to a 3D position point cloud input in the TUM dataset [29]. We observe a similar training behavior, but the test performance diverges.

B. Multiple trajectories learning

To test the relocalisation performance with respect to the number of training sequences, and hence its capability for compression, we use the 7-Scenes dataset ([32], [35]), as shown in Figure 11; this dataset consists of several trajectories taken by different persons moving around the same environment. Training, validation and testing sequences are indicated in the dataset itself.

We first tested this implementation with RGB data and compared it with PoseNet results, and then we evaluate this architecture for all data inputs; results are shown in Table III.

(a) (b) (c) (d)

Fig. 11: Typical views from the Red Kitchen sequence in the 7-Scenes dataset [32].

Although RGB error is again the lowest, with more training images per set, similar performances are found between different data types. However, for the rest of the section we will use only RGB data.

In this test, using RGB data as the sensor input, we observe that CNN-F (8 layers) behaves at least as good as PoseNet, a 23 layers and more complex CNN, where both of them have a relocalisation mean error of 0.59 meters. It remains an open problem the task of designing customized CNN map representation by systematically modifying the neural network architecture itself.

Furthermore, in traditional mapping techniques, in general the map increases when new views are added; instead, by using a CNN map representation, we increasing the map information while maintaining a constant neural network size by re-training it when new trajectories are added.

We used the Red Kitchen, Office and Pumpkin sequences in the 7-Scenes dataset. We took out one of the trajectories for testing and gradually train the CNN adding one remaining sequence at a time. Figure 12 shows that while increasing

TABLE III: CNN-F relocalisation mean error [in meters] for different inputs after 100 epochs in the Red Kitchen sequence from the 7-Scenes dataset[32]. In italics we present the PoseNet mean error reported in [18].

Input	Relocalisation mean error [m]
Depth	1.326 ± 0.372
Gray	0.795 ± 0.504
Point Cloud	0.791 ± 0.692
RGB	**0.589 ± 0.425**
PoseNet (RGB)	*0.59*
RGB+Depth	0.748 ± 0.551
RGB+Point Cloud	0.685 ± 0.682

the number of trajectories, precision also increases but, by construction, the size of the CNN remains the same, as expected. Nevertheless, we didn't reach an asymptotic behavior that would have indicated that the neural network was saturated; this indicates that more trajectories can still be added, improving the performance. While compact, then, this map representation is also constant-size when new information is added.

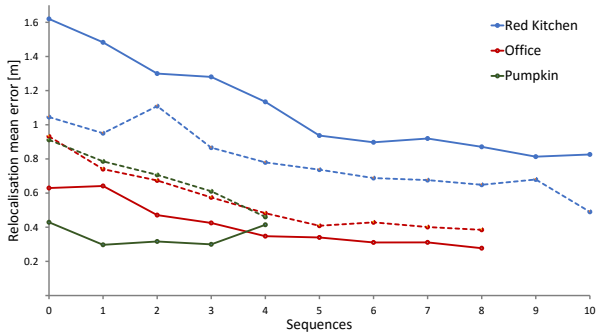

Fig. 12: Relocalisation mean error in the 7-Scenes dataset [32] with respect to the number of training trajectories (we fix one sequence for testing and use the remaining for training; dotted lines indicates a different test sequence). While the number of training trajectories increases, the error decreases but the CNN-F size remains the same (we only affect the weights).

In Figure 8 we can see some outputs for the Red Kitchen sequence where the relocalisation improves as we add more trajectories; in Figure 9 we observe the same behavior in the Office sequence. There, we observe how the relocalised cameras (blue points) are closer to the test sequence (green line) when more training sequences are present (red line).

VI. CONCLUSIONS

We presented a first approach toward CNN map compression for camera relocalisation. We studied the response to different inputs and to different trajectories. We first show that for these kind of models, the RGB images present the best performance, compared with other types of data, as depth or 3D point clouds. We introduced the idea of map compression as the task of finding optimal CNN architectures – we obtain similar performance using CNN-F (8 layers) than PoseNet (23 layers). Then, we demonstrate that increasing the training trajectories, accuracy increases as well, without increasing the map size – i.e. only the filters' weights change but not the number of layers.

For future work we note that more complex relocalisation such as semantic or topological relocalisation were not explored here. One potential direction encouraged by these results is to train simpler networks for object recognition with labeled datasets, and use a second network that accepts semantic information as input for relocalisation. This kind of multi-network systems, where two complex systems interact to perform a single task are of interest to expand on the current work.

REFERENCES

[1] L. Contreras and W. Mayol-Cuevas, "Trajectory-driven point cloud compression techniques for visual SLAM.," in *IROS*, IEEE, 2015.

[2] L. Contreras and W. Mayol-Cuevas, "O-POCO: Online POint cloud COmpression mapping for visual odometry and SLAM.," in *ICRA*, IEEE, 2017 (accepted).

[3] R. Harvey, P. DiCaprio, and K. Heinemann, "A neural network architecture for general image recognition," *Lincoln Laboratory Journal*, vol. 4(2), p. 189207, 1991.

[4] M. J. Milford, G. F. Wyeth, and D. Prasser, "RatSLAM: a hippocampal model for simultaneous localization and mapping," IEEE, 2004.

[5] J. J. DiCarlo, D. Zoccolan, and N. C. Rust., "How does the brain solve visual object recognition?," *Neuron*, vol. 73(3), p. 415434, 2012.

[6] Y. Bengio, A. C. Courville, and P. Vincent, "Unsupervised feature learning and deep learning: A review and new perspectives," *CoRR, abs/1206.5538*, vol. 1, 2012.

[7] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, 2013.

[8] G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.

[9] K. Chatfield, V. Lempitsky, A. Vedaldi, and A. Zisserman, "The devil is in the details: an evaluation of recent feature encoding methods," in *British Machine Vision Conference*, 2011.

[10] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional nets," in *British Machine Vision Conference*, 2014.

[11] P. Agrawal, J. Carreira, and J. Malik, "Learning to see by moving," in *ICCV*, 2015.

[12] K. Jarrett, K. Kavukcuoglu, M. Ranzato, and Y. LeCun, "What is the best multi-stage architecture for object recognition?," in *Proc. International Conference on Computer Vision (ICCV'09)*, IEEE, 2009.

[13] R. Garg, V. K. B. G, and I. D. Reid, "Unsupervised CNN for single view depth estimation: Geometry to the rescue," *CoRR*, vol. abs/1603.04992, 2016.

[14] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *CoRR*, vol. abs/1406.2283, 2014.

[15] B. Li, C. Shen, Y. Dai, A. van den Hengel, and M. He, "Depth and surface normal estimation from monocular images using regression on deep features and hierarchical crfs," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2015.

[16] F. Liu, C. Shen, G. Lin, and I. D. Reid, "Learning depth from single monocular images using deep convolutional neural fields," *CoRR*, vol. abs/1502.07411, 2015.

[17] A. P. Gee and W. Mayol-Cuevas, "6d relocalisation for rgbd cameras using synthetic view regression," in *Proceedings of the British Machine Vision Conference (BMVC)*, September 2012.

[18] A. Kendall, M. Grimes, and R. Cipolla, "Posenet: A convolutional network for real-time 6-dof camera relocalization," 2015.

[19] G. Long, L. Kneip, J. M. Alvarez, and H. Li, "Learning image matching by simply watching video," *CoRR*, vol. abs/1603.06041, 2016.

[20] S. Srinivas, R. K. Sarvadevabhatla, K. R. Mopuri, N. Prabhu, S. S. S. Kruthiventi, and R. V. Babu, "A taxonomy of deep convolutional neural nets for computer vision," *Frontiers in Robotics and AI*, vol. 2, p. 36, 2016.

[21] M. A. Nielsen, *Neural Networks and Deep Learning*. Determination Press, 2015.

[22] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," *CoRR*, vol. abs/1409.4842, 2014.

[23] A. Kendall, M. Grimes, and R. Cipolla, "Research data supporting 'PoseNet: A convolutional network for real-time 6-dof camera relocalization': St marys church [dataset]," 2015. https://www.repository.cam.ac.uk/handle/1810/251294 (accessed February 2016).

[24] M. Cummins and P. Newman, "FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearance," *The International Journal of Robotics Research*, vol. 27, no. 6, pp. 647–665, 2008.

[25] A. Angeli, D. Filliat, S. Doncieux, and J.-A. Meyer, "Incremental vision-based topological slam," in *Proceedings IROS 2008*, 2008.

[26] J. Lim, J. Frahm, and M. Pollefeys, "Online environment mapping using metric-topological maps," *I. J. Robotics Res.*, vol. 31, no. 12, pp. 1394–1408, 2012.

[27] M. Milford and G. Wyeth, "Seqslam : visual route-based navigation for sunny summer days and stormy winter nights," in *Proceedings ICRA 2012*, pp. 1643–1649, IEEE, 2012.

[28] G. James, T. Hastie, R. Tibshirani, and D. Witten, *An Introduction to Statistical Learning: With Applications in R*. Springer, second ed., 2014.

[29] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems.," in *International Conference on Intelligent Robots and Systems*, pp. 573–580, IEEE, 2012.

[30] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "Cnn features off-the-shelf: An astounding baseline for recognition," in *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 512–519, IEEE Computer Society, 2014.

[31] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *CVPR09*, 2009.

[32] B. Glocker, S. Izadi, J. Shotton, and A. Criminisi, "Real-time rgb-d camera relocalization," in *International Symposium on Mixed and Augmented Reality (ISMAR)*, IEEE, 2013.

[33] D. F. Fouhey, A. Gupta, and M. Hebert, "Data-driven 3D primitives for single image understanding," in *ICCV*, 2013.

[34] D. J. Rezende, S. M. A. Eslami, S. Mohamed, P. Battaglia, M. Jaderberg, and N. Heess, "Unsupervised learning of 3d structure from images," *CoRR*, vol. abs/1607.00662, 2016.

[35] J. Shotton, B. Glocker, C. Zach, S. Izadi, A. Criminisi, and A. Fitzgibbon, "Scene coordinate regression forests for camera relocalization in rgb-d images," in *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2013.